

A Pseudo-Euclidean Iteration for Optimal Recovery in Noisy ICA

James Voss*
vossj@cse.ohio-state.edu

Mikhail Belkin*
mbelkin@cse.ohio-state.edu

Luis Rademacher*
lrademac@cse.ohio-state.edu

October 2, 2015

Abstract

Independent Component Analysis (ICA) is a popular model for blind signal separation. The ICA model assumes that a number of independent source signals are linearly mixed to form the observed signals. We propose a new algorithm, PEGI (for pseudo-Euclidean Gradient Iteration), for provable model recovery for ICA with Gaussian noise. The main technical innovation of the algorithm is to use a fixed point iteration in a pseudo-Euclidean (indefinite “inner product”) space. The use of this indefinite “inner product” resolves technical issues common to several existing algorithms for noisy ICA. This leads to an algorithm which is conceptually simple, efficient and accurate in testing.

Our second contribution is combining PEGI with the analysis of objectives for optimal recovery in the noisy ICA model. It has been observed that the direct approach of demixing with the inverse of the mixing matrix is suboptimal for signal recovery in terms of the natural Signal to Interference plus Noise Ratio (SINR) criterion. There have been several partial solutions proposed in the ICA literature. It turns out that any solution to the mixing matrix reconstruction problem can be used to construct an SINR-optimal ICA demixing, despite the fact that SINR itself cannot be computed from data. That allows us to obtain a practical and provably SINR-optimal recovery method for ICA with arbitrary Gaussian noise.

1 Introduction

Independent Component Analysis refers to a class of methods aiming at recovering statistically independent signals by observing their unknown linear combination. There is an extensive literature on this and a number of related problems [7].

Specifically, in the ICA model, we observe n -dimensional realizations $\mathbf{x}(1), \dots, \mathbf{x}(N)$ of a latent variable model $\mathbf{X} = \sum_{k=1}^m S_k A_k = \mathbf{A}\mathbf{S}$ where A_k denotes the k^{th} column of the $n \times m$ mixing matrix A and $\mathbf{S} = (S_1, \dots, S_m)^T$ is the unseen latent random vector of “signals”. It is assumed that S_1, \dots, S_m are independent and non-Gaussian. The source signals and entries of A may be either real- or complex-valued. For simplicity, we will assume throughout that \mathbf{S} has zero mean, as this may be achieved in practice by centering the observed data.

Many ICA algorithms use the preprocessing “whitening” step whose goal is to orthogonalize the independent components. In the noiseless, case this is commonly done by computing the square root of the covariance matrix of \mathbf{X} . Consider now the noisy ICA model $\mathbf{X} = \mathbf{A}\mathbf{S} + \boldsymbol{\eta}$ with additive $\mathbf{0}$ -mean noise $\boldsymbol{\eta}$ independent of \mathbf{S} . It turns out that the introduction of noise makes accurate recovery of the signals significantly more involved. Specifically, whitening using the covariance matrix does not work in the noisy ICA model as the covariance matrix combines both signal and noise. For

*Department of Computer Science and Engineering, the Ohio State University

the case when the noise is Gaussian, matrices constructed from higher order statistics (specifically, cumulants) can be used instead of the covariance matrix. However, these matrices are not in general positive definite and thus the square root cannot always be extracted. This limits the applicability of several previous methods, such as [2, 9, 1]. The GI-ICA algorithm proposed in [21] addresses this issue by using a complicated quasi-orthogonalization step followed by an iterative method.

In this paper (section 2), we develop a simple and practical one-step algorithm, PEGI (for pseudo-Euclidean Gradient Iteration) for provably recovering A (up to the unavoidable ambiguities of the model) in the case when the noise is Gaussian (with an arbitrary, unknown covariance matrix). The main technical innovation of our approach is to formulate the recovery problem as a fixed point method in an indefinite “inner product” (pseudo-Euclidean) space.

The second contribution of the paper is combining PEGI with the analysis of objectives for optimal recovery in the noisy ICA model. In most applications of ICA (e.g., speech separation [18], MEG/EEG artifact removal [20] and others) one cares about recovering the signals $\mathbf{s}(1), \dots, \mathbf{s}(N)$. This is known as the *source recovery problem*. This is typically done by first recovering the matrix A (up to an appropriate scaling of the column directions).

At first, source recovery and recovering the mixing matrix A appear to be essentially equivalent. In the *noiseless* ICA model, if A is invertible¹ then $\mathbf{s}(t) = A^{-1}\mathbf{x}(t)$ recovers the sources. On the other hand, in the *noisy* model, the exact recovery of the latent sources $\mathbf{s}(t)$ becomes impossible even if A is known exactly. Part of the “noise” can be incorporated into the “signal” preserving the form of the model. Even worse, neither A nor \mathbf{S} are defined uniquely as there is an inherent ambiguity in the setting. There could be many equivalent decompositions of the observed signal as $\mathbf{X} = A'\mathbf{S}' + \boldsymbol{\eta}'$ (see the discussion in section 3).

We consider recovered signals of the form $\hat{\mathbf{S}}(B) := B\mathbf{X}$ for a choice of $m \times n$ demixing matrix B . Signal recovery is considered optimal if the coordinates of $\hat{\mathbf{S}}(B) = (\hat{S}_1(B), \dots, \hat{S}_m(B))$ maximize Signal to Interference plus Noise Ratio (SINR) within any *fixed* model $\mathbf{X} = A\mathbf{S} + \boldsymbol{\eta}$. Note that the value of SINR depends on the decomposition of the observed data into “noise” and “signal”: $\mathbf{X} = A'\mathbf{S}' + \boldsymbol{\eta}'$.

Surprisingly, the SINR optimal demixing matrix does not depend on the decomposition of data into signal plus noise. As such, SINR optimal ICA recovery is well defined given access to data despite the inherent ambiguity in the model. Further, it will be seen that the SINR optimal demixing can be constructed from $\text{cov}(\mathbf{X})$ and the directions of the columns of A (which are also well-defined across signal/noise decompositions).

Our SINR-optimal demixing approach combined with the PEGI algorithm provides a complete SINR-optimal recovery algorithm in the ICA model with arbitrary Gaussian noise. We note that the ICA papers of which we are aware that discuss optimal demixing do not observe that SINR optimal demixing is invariant to the choice of signal/noise decomposition. Instead, they propose more limited strategies for improving the demixing quality within a fixed ICA model. For instance, Joho et al. [14] show how SINR-optimal demixing can be approximated with extra sensors when assuming a white additive noise, and Koldovský and Tichavský [16] discuss how to achieve asymptotically low bias ICA demixing assuming white noise within a fixed ICA model. However, the invariance of the SINR-optimal demixing matrix appears in the array sensor systems literature [6].

Finally, in section 4, we demonstrate experimentally that our proposed algorithm for ICA outperforms existing practical algorithms at the task of noisy signal recovery, including those specifically designed for beamforming, when given sufficiently many samples. Moreover, most existing practical algorithms for noisy source recovery have a bias and cannot recover the optimal demixing matrix even with infinite samples. We also show that PEGI requires significantly fewer samples

¹ A^{-1} can be replaced with A^\dagger in the discussion below for over-determined ICA.

than GI-ICA [21] to perform ICA accurately.

1.1 The Indeterminacies of ICA

Notation: We use M^* to denote the entry-wise complex conjugate of a matrix M , M^T to denote its transpose, and M^H to denote its conjugate transpose.

Before proceeding with our results, we discuss the somewhat subtle issue of indeterminacies in ICA. These ambiguities arise from the fact that the observed \mathbf{X} may have multiple decompositions into ICA models $\mathbf{X} = \mathbf{A}\mathbf{S} + \boldsymbol{\eta}$ and $\mathbf{X} = \mathbf{A}'\mathbf{S}' + \boldsymbol{\eta}'$.

Noise-free ICA has two natural indeterminacies. For any nonzero constant α , the contribution of the k^{th} component $A_k S_k$ to the model can equivalently be obtained by replacing A_k with αA_k and S_k with the rescaled signal $\frac{1}{\alpha} S_k$. To lessen this scaling indeterminacy, we use the convention² that $\text{cov}(\mathbf{S}) = \mathcal{I}$ throughout this paper. As such, each source S_k (or equivalently each A_k) is defined up to a choice of sign (a unit modulus factor in the complex case). In addition, there is an ambiguity in the order of the latent signals. For any permutation π of $[m]$ (where $[m] := \{1, \dots, m\}$), the ICA models $\mathbf{X} = \sum_{k=1}^m S_k A_k$ and $\mathbf{X} = \sum_{k=1}^m S_{\pi(k)} A_{\pi(k)}$ are indistinguishable. In the noise free setting, \mathbf{A} is said to be recovered if we recover each column of \mathbf{A} up to a choice of sign (or up to a unit modulus factor in the complex case) and an unknown permutation. As the sources S_1, \dots, S_m are only defined up to the same indeterminacies, inverting the recovered matrix $\tilde{\mathbf{A}}$ to obtain a demixing matrix works for signal recovery.

In the noisy ICA setting, there is an additional indeterminacy in the definition of the sources. Consider $\boldsymbol{\xi}$ to be $\mathbf{0}$ -mean axis-aligned Gaussian random vector. Then, the noisy ICA model $\mathbf{X} = \mathbf{A}(\mathbf{S} + \boldsymbol{\xi}) + \boldsymbol{\eta}$ in which $\boldsymbol{\xi}$ is considered part of the latent source signal $\mathbf{S}' = \mathbf{S} + \boldsymbol{\xi}$, and the model $\mathbf{X} = \mathbf{A}\mathbf{S} + (\mathbf{A}\boldsymbol{\xi} + \boldsymbol{\eta})$ in which $\boldsymbol{\xi}$ is part of the noise are indistinguishable. In particular, the latent source \mathbf{S} and its covariance are ill-defined. Due to this extra indeterminacy, the lengths of the columns of \mathbf{A} no longer have a fully defined meaning even when we assume $\text{cov}(\mathbf{S}) = \mathcal{I}$. In the noisy setting, \mathbf{A} is said to be recovered if we obtain the columns of \mathbf{A} up to non-zero scalar multiplicative factors and an arbitrary permutation.

The last indeterminacy is the most troubling as it suggests that the power of each source signal is itself ill-defined in the noisy setting. Despite this indeterminacy, it is possible to perform an SINR-optimal demixing without additional assumptions about what portion of the signal is source and what portion is noise. In section 3, we will see that SINR-optimal source recovery takes on a simple form: Given any solution $\tilde{\mathbf{A}}$ which recovers \mathbf{A} up to the inherent ambiguities of noisy ICA, then $\tilde{\mathbf{A}}^H \text{cov}(\mathbf{X})^\dagger$ is an SINR-optimal demixing matrix.

1.2 Related Work and Contributions

Independent Component Analysis is probably the most used model for Blind Signal Separation. It has seen numerous applications and has generated a vast literature, including in the noisy and underdetermined settings. We refer the reader to the books [13, 7] for a broad overview of the subject.

It was observed early on by Cardoso [4] that ICA algorithms based solely on higher order cumulant statistics are invariant to additive Gaussian noise. This observation has allowed the creation of many algorithms for recovering the ICA mixing matrix in the noisy and often underdetermined settings. Despite the significant work on noisy ICA algorithms, they remain less efficient, more specialized, or less practical than the most popular noise free ICA algorithms.

²Alternatively, one may place the scaling information in the signals by setting $\|A_k\| = 1$ for each k .

Research on cumulant-based noisy ICA can largely be split into several lines of work which we only highlight here. Some algorithms such as FOObI [4] and BIOME [1] directly use the tensor structure of higher order cumulants. In another line of work, De Lathauwer et al. [8] and Yeredor [23] have suggested algorithms which jointly diagonalize cumulant matrices in a manner reminiscent of the noise-free JADE algorithm [3]. In addition, Yeredor [22] and Goyal et al. [11] have proposed ICA algorithms based on random directional derivatives of the second characteristic function.

Each line of work has its advantages and disadvantages. The joint diagonalization algorithms and the tensor based algorithms tend to be practical in the sense that they use redundant cumulant information in order to achieve more accurate results. However, they have a higher memory complexity than popular noise free ICA algorithms such as FastICA [12]. Moreover, the tensor methods (FOObI and BIOME) require the latent source signals to have positive order $2k$ ($k \geq 2$, a predetermined fixed integer) cumulants as they rely on taking a matrix square root. Finally, the methods based on random directional derivatives of the second characteristic function rely heavily upon randomness in a manner not required by the most popular noise free ICA algorithms.

We continue a line of research started by Arora et al. [2] and Voss et al. [21] on fully determined noisy ICA which addresses some of these practical issues by using a deflationary approach reminiscent of FastICA. Their algorithms thus have lower memory complexity and are more scalable to high dimensional data than the joint diagonalization and tensor methods. However, both works require a preprocessing step (quasi-orthogonalization) to orthogonalize the latent signals which is based on taking a matrix square root. Arora et al. [2] require each latent signal to have positive fourth cumulant in order to carry out this preprocessing step. In contrast, Voss et al. [21] are able to perform quasi-orthogonalization with source signals of mixed sign fourth cumulants; but their quaise-orthogonalization step is more complicated and can run into numerical issues under sampling error. We demonstrate that quasi-orthogonalization is unnecessary. We introduce the PEGI algorithm to work within a (not necessarily positive definite) inner product space instead. Experimentally, this leads to improved demixing performance. In addition, we handle the case of complex signals.

Finally, another line of work attempts to perform SINR-optimal source recovery in the noisy ICA setting. It was noted by Koldovský and Tichavský [15] that for noisy ICA, traditional ICA algorithms such as FastICA and JADE actually outperform algorithms which first recover A in the noisy setting and then use the resulting approximation of A^\dagger to perform demixing. It was further observed that A^\dagger is not the optimal demixing matrix for source recovery. Later, Koldovský and Tichavský [17] proposed an algorithm based on FastICA which performs a low SINR-bias beamforming.

2 Pseudo-Euclidean Gradient Iteration ICA

In this section, we introduce the PEGI algorithm for recovering A in the “fully determined” noisy ICA setting where $m \leq n$. PEGI relies on the idea of Gradient Iteration introduced Voss et al. [21]. However, unlike GI-ICA Voss et al. [21], PEGI does not require the source signals to be orthogonalized. As such, PEGI does not require the complicated quasi-orthogonalization preprocessing step of GI-ICA which can be inaccurate to compute in practice. We sketch the Gradient Iteration algorithm in Section 2.1, and then introduce PEGI in Section 2.2. For simplicity, we limit this discussion to the case of real-valued signals. We show how to construct PEGI for complex-valued signals in Appendix A.

In this section we assume a noisy ICA model $\mathbf{X} = \mathbf{A}\mathbf{S} + \boldsymbol{\eta}$ such that $\boldsymbol{\eta}$ is arbitrary Gaussian and independent of \mathbf{S} . We also assume that $m \leq n$, that m is known, and that the columns of A

are linearly independent.

2.1 Gradient Iteration with Orthogonality

The gradient iteration relies on the properties of cumulants. We will focus on the fourth cumulant, though similar constructions may be given using other even order cumulants of higher order. For a zero-mean random variable X , the fourth order cumulant may be defined as $\kappa_4(X) := \mathbb{E}[X^4] - 3\mathbb{E}[X^2]^2$ [see 7, Chapter 5, Section 1.2]. Higher order cumulants have nice algebraic properties which make them useful for ICA. In particular, κ_4 has the following properties: (1) (Independence) If X and Y are independent, then $\kappa_4(X + Y) = \kappa_4(X) + \kappa_4(Y)$. (2) (Homogeneity) If α is a scalar, then $\kappa_4(\alpha X) = \alpha^4 \kappa_4(X)$. (3) (Vanishing Gaussians) If X is normally distributed then $\kappa_4(X) = 0$.

We consider the following function defined on the unit sphere: $f(\mathbf{u}) := \kappa_4(\langle \mathbf{X}, \mathbf{u} \rangle)$. Expanding $f(\mathbf{u})$ using the above properties we obtain:

$$f(\mathbf{u}) = \kappa_4\left(\sum_{k=1}^m \langle A_k, \mathbf{u} \rangle S_k + \langle \mathbf{u}, \boldsymbol{\eta} \rangle\right) = \sum_{k=1}^m \langle A_k, \mathbf{u} \rangle^4 \kappa_4(S_k).$$

Taking derivatives we obtain:

$$\nabla f(\mathbf{u}) = 4 \sum_{k=1}^m \langle A_k, \mathbf{u} \rangle^3 \kappa_4(S_k) A_k \quad (1)$$

$$\mathcal{H}f(\mathbf{u}) = 12 \sum_{k=1}^m \langle A_k, \mathbf{u} \rangle^2 \kappa_4(S_k) A_k A_k^T = AD(\mathbf{u})A^T \quad (2)$$

where $D(\mathbf{u})$ is a diagonal matrix with entries $D(\mathbf{u})_{kk} = 12 \langle A_k, \mathbf{u} \rangle^2 \kappa_4(S_k)$.

Voss et al. [21] introduced GI-ICA as a fixed point algorithm under the assumption that the columns of A are orthogonal but not necessarily unit vectors. The main idea is that the update $\mathbf{u} \leftarrow \nabla f(\mathbf{u}) / \|\nabla f(\mathbf{u})\|$ is a form of a generalized power iteration. From equation (1), each A_k may be considered as a direction in a hidden orthogonal basis of the space. During each iteration, the A_k coordinate of \mathbf{u} is raised to the 3rd power and multiplied by a constant. Treating this iteration as a fixed point update, it was shown that given a random starting point, this iterative procedure converges rapidly to one of the columns of A (up to a choice of sign). The rate of convergence is cubic.

However, the GI-ICA algorithm requires a somewhat complicated preprocessing step called quasi-orthogonalization to linearly transform the data to make columns of A orthogonal. Quasi-orthogonalization makes use of evaluations of Hessians of the fourth cumulant function to construct a matrix of the form $C = ADA^T$ where D has all positive diagonal entries—a task which is complicated by the possibility that the latent signals S_i may have fourth order cumulants of differing signs—and requires taking the matrix square root of a positive definite matrix of this form. However, the algorithm used for constructing C under sampling error is not always positive definite in practice, which can make the preprocessing step fail. We will show how our PEGI algorithm makes quasi-orthogonalization unnecessary, in particular, resolving this issue.

2.2 Gradient Iteration in a Pseudo-Euclidean Space

We now show that the gradient iteration can be performed using in a pseudo-Euclidean space in which the columns of A are orthogonal. The natural candidate for the “inner product space” would be to use $\langle \cdot, \cdot \rangle_*$ defined as $\langle \mathbf{u}, \mathbf{v} \rangle_* := \mathbf{u}^T (AA^T)^\dagger \mathbf{v}$. Clearly, $\langle A_i, A_j \rangle_* = \delta_{ij}$ gives the desired

Algorithm 1 Recovers a column of A up to a scaling factor if \mathbf{u}_0 is generically chosen.

Inputs: Unit vector \mathbf{u}_0 , C , ∇f
 $k \leftarrow 1$
repeat
 $\mathbf{u}_k \leftarrow \nabla f(C^\dagger \mathbf{u}_{k-1}) / \|\nabla f(C^\dagger \mathbf{u}_{k-1})\|$
 $k \leftarrow k + 1$
until Convergence (up to sign)
return \mathbf{u}_k

orthogonality property. However, there are two issues with this “inner product space”: First, it is only an inner product space when A is invertible. This turns out not to be a major issue, and we move forward largely ignoring this point. The second issue is more fundamental: We only have access to AA^T in the noise free setting where $\text{cov}(\mathbf{X}) = AA^T$. In the noisy setting, we have access to matrices of the form $\mathcal{H}f(\mathbf{u}) = AD(\mathbf{u})A^T$ from equation (2) instead. consider a pseudo-Euclidean inner product defined as follows: Let $C = ADA^T$ where D is a diagonal matrix with non-zero diagonal entries, and define $\langle \cdot, \cdot \rangle_C$ by $\langle \mathbf{u}, \mathbf{v} \rangle_C = \mathbf{u}^T C^\dagger \mathbf{v}$. When D contains negative entries, this is not a proper inner product since C is not positive definite. In particular, $\langle A_k, A_k \rangle_C = A_k^T (ADA^T)^\dagger A_k = d_{kk}^{-1}$ may be negative. Nevertheless, when $k \neq j$, $\langle A_k, A_j \rangle_C = A_k^T (ADA^T)^\dagger A_j = 0$ gives that the columns of A are orthogonal in this space.

We define functions $\alpha_k : \mathbb{R}^n \rightarrow \mathbb{R}$ by $\alpha_k(\mathbf{u}) = (A^\dagger \mathbf{u})_k$ such that for any $\mathbf{u} \in \text{span}(A_1, \dots, A_m)$, then $\mathbf{u} = \sum_{i=1}^m \alpha_i(\mathbf{u}) A_i$ is the expansion of \mathbf{u} in its A_i basis. Continuing from equation (1), for any $\mathbf{u} \in S^{n-1}$ we see $\nabla f(C^\dagger \mathbf{u}) = 4 \sum_{k=1}^n \langle A_k, C^\dagger \mathbf{u} \rangle^3 \kappa_4(S_k) A_k = 4 \sum_{k=1}^n \langle A_k, \mathbf{u} \rangle_C^3 \kappa_4(S_k) A_k$ is the gradient iteration recast in the $\langle \cdot, \cdot \rangle_C$ space. Expanding \mathbf{u} in its A_k basis, we obtain

$$\nabla f(C^\dagger \mathbf{u}) = 4 \sum_{k=1}^m (\alpha_k(\mathbf{u}) \langle A_k, A_k \rangle_C)^3 \kappa_4(S_k) A_k = 4 \sum_{k=1}^m \alpha_k(\mathbf{u})^3 (d_{kk}^{-3} \kappa_4(S_k)) A_k, \quad (3)$$

which is a power iteration in the unseen A_k coordinate system. As no assumptions are made upon the $\kappa_4(S_k)$ values, the d_{kk}^{-3} scalings which were not present in eq. (1) cause no issues. Using this update, we obtain Alg. 1, a fixed point method for recovering a single column of A up to an unknown scaling.

Before proceeding, we should clarify the notion of fixed point convergence in Algorithm 1. We say that the sequence $\{\mathbf{u}_k\}_{k=0}^\infty$ converges to \mathbf{v} up to sign if there exists a sequence $\{c_k\}_{k=0}^\infty$ such that each $c_k \in \{\pm 1\}$ and $c_k \mathbf{u}_k \rightarrow \mathbf{v}$ as $k \rightarrow \infty$. We have the following convergence guarantee.

Theorem 1. *If \mathbf{u}_0 is chosen uniformly at random from S^{n-1} , then with probability 1, there exists $\ell \in [m]$ such that the sequence $\{\mathbf{u}_k\}_{k=0}^\infty$ defined as in Algorithm 1 converges to $A_\ell / \|A_\ell\|$ up to sign. Further, the rate of convergence is cubic.*

Due to limited space, we omit the proof of Theorem 1. It is similar to the proof of [21, Theorem 4].

In practice, we test near convergence by checking if we are still making significant progress. In particular, for some predefined $\epsilon > 0$, if there exists a sign value $c_k \in \{\pm 1\}$ such that $\|\mathbf{u}_k - c_k \mathbf{u}_{k-1}\| < \epsilon$, then we declare convergence achieved and return the result. As there are only two choices for c_k , this is easily checked, and we exit the loop if this condition is met.

Full ICA Recovery Via the Pseudo-Euclidean GI-Update. We are able to recover a single column of A up to its unknown scale. However, for full recovery of A , we would like (given recovered columns $A_{\ell_1}, \dots, A_{\ell_j}$) to be able to recover a column A_k such that $k \notin \{\ell_1, \dots, \ell_j\}$ on demand.

Algorithm 2 Full ICA matrix recovery algorithm. Returns two matrices: (1) \tilde{A} is the recovered mixing matrix for the noisy ICA model $\mathbf{X} = \mathbf{A}\mathbf{S} + \boldsymbol{\eta}$, and (2) \tilde{B} is a running estimate of \tilde{A}^\dagger .

```

1: Inputs:  $C, \nabla f$ 
2:  $\tilde{A} \leftarrow 0, \tilde{B} \leftarrow 0$ 
3: for  $j \leftarrow 1$  to  $m$  do
4:   Draw  $\mathbf{u}$  uniformly at random from  $S^{n-1}$ .
5:   repeat
6:      $\mathbf{u} \leftarrow \mathbf{u} - \tilde{A}\tilde{B}\mathbf{u}$ 
7:      $\mathbf{u} \leftarrow \nabla f(C^\dagger \mathbf{u}) / \|\nabla f(C^\dagger \mathbf{u})\|$ .
8:   until Convergence (up to sign)
9:    $\tilde{A}_j \leftarrow \mathbf{u}$ 
10:   $\tilde{B}_j \leftarrow [C^\dagger A_j / ((C^\dagger A_j)^T A_j)]^T$ 
11: end for
12: return  $\tilde{A}, \tilde{B}$ 

```

The idea behind the simultaneous recovery of all columns of A is two-fold. First, instead of just finding columns of A using Algorithm 1, we simultaneously find rows of A^\dagger . Then, using the recovered columns of A and rows of A^\dagger , we project \mathbf{u} onto the orthogonal complement of the recovered columns of A within the $\langle \cdot, \cdot \rangle_C$ pseudo-inner product space.

Recovering rows of A^\dagger . Suppose we have access to a column A_k (which may be achieved using Algorithm 1). Let $A_{k\cdot}^\dagger$ denote the k^{th} row of A^\dagger . Then, we note that $C^\dagger A_k = (ADA^T)^\dagger A_k = d_{kk}^{-1}(A^T)^\dagger_k = d_{kk}^{-1}(A_{k\cdot}^\dagger)^T$ recovers $A_{k\cdot}^\dagger$ up to an arbitrary, unknown constant d_{kk}^{-1} . However, the constant d_{kk}^{-1} may be recovered by noting that $\langle A_k, A_k \rangle_C = (C^\dagger A_k)^T A_k = d_{kk}^{-1}$. As such, we may estimate $A_{k\cdot}^\dagger$ as $[C^\dagger A_k / ((C^\dagger A_k)^T A_k)]^T$.

Enforcing Orthogonality During the GI Update. Given access to a vector $\mathbf{u} = \sum_{k=1}^m \alpha_k(\mathbf{u})A_k + P_{A^\perp} \mathbf{u}$ (where P_{A^\perp} is the projection onto the orthogonal complements of the range of A), some recovered columns $A_{\ell_1}, \dots, A_{\ell_r}$, and corresponding rows of A^\dagger , we may zero out the components of \mathbf{u} corresponding to the recovered columns of A . Letting $\mathbf{u}' = \mathbf{u} - \sum_{j=1}^r A_{\ell_j} A_{\ell_j}^\dagger \mathbf{u}$, then $\mathbf{u}' = \sum_{k \in [m] \setminus \{\ell_1, \dots, \ell_r\}} \alpha_k(\mathbf{u})A_k + P_{A^\perp} \mathbf{u}$. In particular, \mathbf{u}' is orthogonal (in the $\langle \cdot, \cdot \rangle_C$ space) to the previously recovered columns of A . This allows the non-orthogonal gradient iteration algorithm to recover a new column of A .

Using these ideas, we obtain Algorithm 2, which is the PEGI algorithm for recovery of the mixing matrix A in noisy ICA up to the inherent ambiguities of the problem. Within this Algorithm, step 6 enforces orthogonality with previously found columns of A , guaranteeing that convergence to a new column of A .

Practical Construction of C . In our implementation, we set $C = \frac{1}{12} \sum_{k=1}^n \mathcal{H}f(\mathbf{e}_k)$, as it can be shown from equation (2) that $\sum_{k=1}^n \mathcal{H}f(\mathbf{e}_k) = ADA^T$ with $d_{kk} = \|A_k\|^2 \kappa_4(S_k)$. This deterministically guarantees that each latent signal has a significant contribution to C .

3 SINR Optimal Recovery in Noisy ICA

In this section, we demonstrate how to perform SINR optimal ICA within the noisy ICA framework given access to an algorithm (such as PEGI) to recover the directions of the columns of A . To this end, we first discuss the SINR optimal demixing solution within any decomposition of the ICA

model into signal and noise as $\mathbf{X} = \mathbf{A}\mathbf{S} + \boldsymbol{\eta}$. We then demonstrate that the SINR optimal demixing matrix is actually the same across all possible model decompositions, and that it can be recovered. The results in this section hold in greater generality than in section 2. They hold even if $m \geq n$ (the underdetermined setting) and even if the additive noise $\boldsymbol{\eta}$ is non-Gaussian.

Consider B an $m \times n$ demixing matrix, and define $\hat{\mathbf{S}}(B) := B\mathbf{X}$ the resulting approximation to \mathbf{S} . It will also be convenient to estimate the source signal \mathbf{S} one coordinate at a time: Given a row vector \mathbf{b} , we define $\hat{S}(\mathbf{b}) := \mathbf{b}\mathbf{X}$. If $\mathbf{b} = B_{k\cdot}$ (the k^{th} row of B), then $\hat{S}(\mathbf{b}) = [\hat{\mathbf{S}}(B)]_k = \hat{S}_k(B)$ is our estimate to the k^{th} latent signal S_k . Within a specific ICA model $\mathbf{X} = \mathbf{A}\mathbf{S} + \boldsymbol{\eta}$, signal to interference-plus-noise ratio (SINR) is defined by the following equation:

$$\text{SINR}_k(\mathbf{b}) := \frac{\text{var}(\mathbf{b}A_k S_k)}{\text{var}(\mathbf{b}\mathbf{A}\mathbf{S} - \mathbf{b}A_k S_k) + \text{var}(\mathbf{b}\boldsymbol{\eta})} = \frac{\text{var}(\mathbf{b}A_k S_k)}{\text{var}(\mathbf{b}\mathbf{A}\mathbf{X}) - \text{var}(\mathbf{b}A_k S_k)}. \quad (4)$$

SINR_k is the variance of the contribution of k^{th} source divided by the variance of the noise and interference contributions within the signal.

Given access to the mixing matrix A , we define $B_{\text{opt}} = A^H(AA^H + \text{cov}(\boldsymbol{\eta}))^\dagger$. Since $\text{cov}(\mathbf{X}) = AA^H + \text{cov}(\boldsymbol{\eta})$, this may be rewritten as $B_{\text{opt}} = A^H \text{cov}(\mathbf{X})^\dagger$. Here, $\text{cov}(\mathbf{X})^\dagger$ may be estimated from data, but due to the ambiguities of the noisy ICA model, A (and specifically its column norms) cannot be estimated from data.

Koldovský and Tichavský [15] observed that when $\boldsymbol{\eta}$ is a white Gaussian noise, B_{opt} jointly maximizes SINR_k for each $k \in [m]$, i.e., SINR_k takes on its maximal value at $(B_{\text{opt}})_{k\cdot}$. Below in Proposition 2, we generalize this result to include arbitrary non-spherical, potentially non-Gaussian noise.

It is interesting to note that even after the data is whitened, i.e. $\text{cov}(X) = \mathcal{I}$, the optimal SINR solution is different from the optimal solution in the noiseless case unless A is an orthogonal matrix, i.e. $A^\dagger = A^H$. This is generally not the case, even if $\boldsymbol{\eta}$ is white Gaussian noise.

Proposition 2. *For each $k \in [m]$, $(B_{\text{opt}})_{k\cdot}$ is a maximizer of SINR_k .*

The proof of Proposition 2 is deferred to appendix B.

Since SINR is scale invariant, Proposition 2 implies that any matrix of the form $DB_{\text{opt}} = DA^H \text{cov}(\mathbf{X})^\dagger$ where D is a diagonal scaling matrix (with non-zero diagonal entries) is an SINR-optimal demixing matrix. More formally, we have the following result.

Theorem 3. *Let \tilde{A} be an $n \times m$ matrix containing the columns of A up to scale and an arbitrary permutation. That is, there exists a permutation π of $[m]$ and non-zero constants $\alpha_1, \dots, \alpha_m$ such that $\alpha_k \tilde{A}_{\pi(k)} = A_k$ for each $k \in [m]$. Then, $(\tilde{A}^H \text{cov}(X)^\dagger)_{\pi(k)\cdot}$ is a maximizer of SINR_k .*

By Theorem 3, given access to a matrix \tilde{A} which recovers the directions of the columns of A , then $\tilde{A}^H \text{cov}(X)^\dagger$ is the SINR-optimal demixing matrix. For ICA in the presence of Gaussian noise, the directions of the columns of A are well defined simply from \mathbf{X} , that is, the directions of the columns of A do not depend on the decomposition of \mathbf{X} into signal and noise (see the discussion in section 1.1 on ICA indeterminacies). The problem of SINR optimal demixing is thus well defined for ICA in the presence of Gaussian noise, and the SINR optimal demixing matrix can be estimated from data without any additional assumptions on the magnitude of the noise in the data.

Finally, we note that in the noise-free case, the SINR-optimal source recovery simplifies to be \tilde{A}^\dagger .

Corollary 4. *Suppose that $\mathbf{X} = \mathbf{A}\mathbf{S}$ is a noise free (possibly underdetermined) ICA model. Suppose that $\tilde{A} \in \mathbb{R}^{n \times m}$ contains the columns of A up to scale and permutation, i.e., there exists diagonal matrix D with non-zero entries and a permutation matrix Π such that $\tilde{A} = A D \Pi$. Then \tilde{A}^\dagger is an SINR-optimal demixing matrix.*

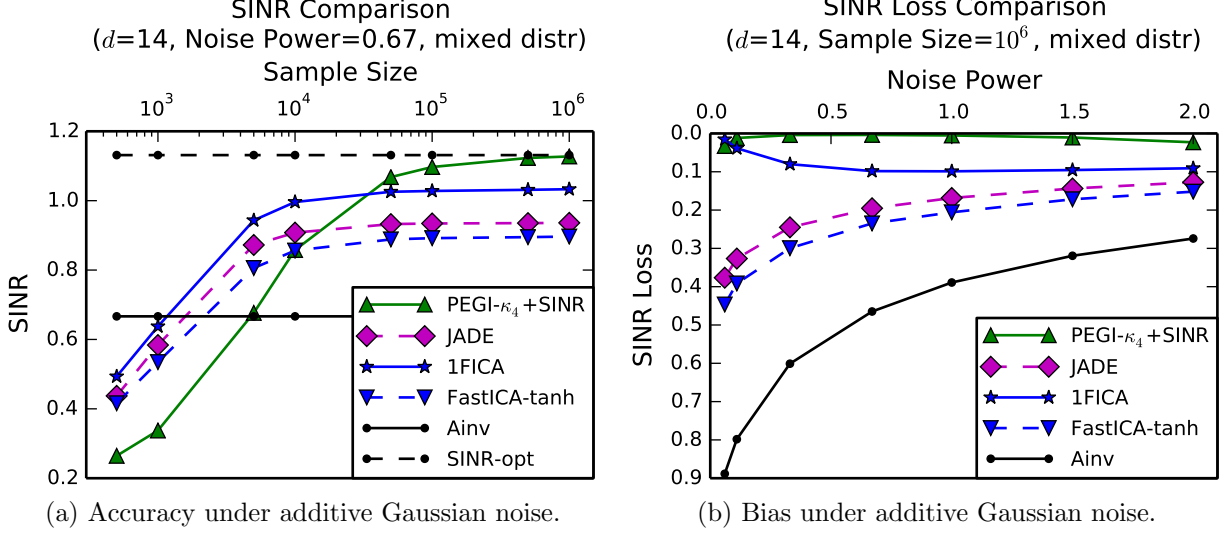


Figure 1: SINR performance comparison of ICA algorithms.

Proof. By Theorem 3, $(AD^{-1}\Pi)^H \text{cov}(\mathbf{X})^\dagger$ is an SINR-optimal demixing matrix. Expanding, we obtain: $(AD^{-1}\Pi)^H \text{cov}(\mathbf{X})^\dagger = \Pi^H D^{-1} A^H (AA^H)^\dagger = \Pi^H D^{-1} A^\dagger = (AD\Pi)^\dagger = \tilde{A}^\dagger$. \square

Corollary 4 is consistent with known beamforming results. In particular, it is known that A^\dagger is an optimal (in terms of minimum mean squared error) beamforming matrix for underdetermined ICA [19, section 3B].

4 Experimental Results

We now compare the proposed PEGI algorithm with several existing ICA algorithms. In addition to qorth+GI-ICA (that is, GI-ICA with the quasi-orthogonalization preprocessing step), we use the following algorithmic baselines:

JADE [3] is a popular fourth cumulant based ICA algorithm designed for the noise free setting. We use the implementation of Cardoso and Souloumiac [5].

FastICA [12] is a popular ICA algorithm designed for the noise free setting based on a deflationary approach of recovering one component at a time. We use the implementation of Gävert et al. [10].

1FICA [16, 17] is a variation of FastICA with the tanh contrast function designed to have low bias for performing SINR-optimal beamforming in the presence of Gaussian noise.

Ainv is the oracle demixing algorithm which uses A^\dagger as the demixing matrix.

SINR-opt is the oracle demixing algorithm which demixes using $A^H \text{cov}(\mathbf{X})^\dagger$ to achieve an SINR-optimal demixing.

We compare these algorithms on simulated data with $n = m$. We constructed mixing matrices A with condition number 3 via a reverse singular value decomposition ($A = U\Lambda V^T$). The matrices U and V were random orthogonal matrices, and Λ was chosen to have 1 as its minimum and 3 as its maximum singular values, with the intermediate singular values chosen uniformly at random. We drew data from a noisy ICA model $\mathbf{X} = \mathbf{AS} + \boldsymbol{\eta}$ where $\text{cov}(\boldsymbol{\eta}) = \Sigma$ was chosen to be maligned with $\text{cov}(\mathbf{AS}) = \mathbf{AA}^T$. We set $\Sigma = p(10\mathbf{I} - \mathbf{AA}^T)$ where p is a constant defining the *noise power*. It can be shown that $p = \frac{\max_{\mathbf{v}} \text{var}(\mathbf{v}^T \boldsymbol{\eta})}{\max_{\mathbf{v}} \text{var}(\mathbf{v}^T \mathbf{AS})}$ is the ratio of the maximum directional noise variance to the maximum directional signal variance. We generated 100 matrices A for our experiments with 100 corresponding ICA data sets for each sample size and noise power. When reporting




Figure 2: Accuracy comparison of PEGI using pseudo-inner product spaces and GI-ICA using quasi-orthogonalization.

results, we apply each algorithm to each of the 100 data sets for the corresponding sample size and noise power and we report the mean performance. The source distributions used in our ICA experiments were the Laplace and Bernoulli distribution with parameters 0.05 and 0.5 respectively, the t -distribution with 3 and 5 degrees of freedom respectively, the exponential distribution, and the uniform distribution. Each distribution was normalized to have unit variance, and the distributions were each used twice to create 14-dimensional data. We compare the algorithms using either SINR or the SINR loss from the optimal demixing matrix (defined by $\text{SINR Loss} = [\text{Optimal SINR} - \text{Achieved SINR}]$).

In Figure 1b, we compare our proposed ICA algorithm with various ICA algorithms for signal recovery. In the PEGI- κ_4 +SINR algorithm, we use PEGI- κ_4 to estimate A , and then perform demixing using the resulting estimate of $A^H \text{cov}(\mathbf{X})^{-1}$, the formula for SINR-optimal demixing. It is apparent that when given sufficient samples, PEGI- κ_4 +SINR provides the best SINR demixing. JADE, FastICA-tanh, and 1FICA each have a bias in the presence of additive Gaussian noise which keeps them from being SINR-optimal even when given many samples.

In Figure 1a, we compare algorithms at various sample sizes. The PEGI- κ_4 +SINR algorithm relies more heavily on accurate estimates of fourth order statistics than JADE, and the FastICA-tanh and 1FICA algorithms do not require the estimation of fourth order statistics. For this reason, PEGI- κ_4 +SINR requires more samples than the other algorithms in order to be run accurately. However, once sufficient samples are taken, PEGI- κ_4 +SINR outperforms the other algorithms including 1FICA which is designed to have low SINR bias.

In order to avoid clutter, we did not include qorth+GI-ICA- κ_4 +SINR (the SINR optimal demixing estimate constructed using qorth+GI-ICA- κ_4 to estimate A) in the figures 1b and 1a. It is also asymptotically unbiased in estimating the directions of the columns of A , and similar conclusions could be drawn using qorth+GI-ICA- κ_4 in place of PEGI- κ_4 . However, in Figure 2, we see that PEGI- κ_4 +SINR requires fewer samples than qorth+GI-ICA- κ_4 +SINR to achieve good performance. This is particularly highlighted in the medium sample regime.

On the Performance of Traditional ICA Algorithms for Noisy ICA. An interesting observation [first made in 15] is that the popular noise free ICA algorithms JADE and FastICA perform reasonably well in the noisy setting. In Figures 1b and 1a, they significantly outperform

demixing using A^{-1} for source recovery. It turns out that this may be explained by a shared preprocessing step. Both JADE and FastICA rely on a whitening preprocessing step in which the data are linearly transformed to have identity covariance. It can be shown in the noise free setting that after whitening, the mixing matrix A is a rotation matrix. These algorithms proceed by recovering an orthogonal matrix \tilde{A} to approximate the true mixing matrix A . Demixing is performed using $\tilde{A}^{-1} = \tilde{A}^H$. Since the data is white (has identity covariance), then the demixing matrix $\tilde{A}^H = \tilde{A}^H \text{cov}(\mathbf{X})^{-1}$ is an estimate of the SINR-optimal demixing matrix. Nevertheless, the traditional ICA algorithms give a biased estimate of A under additive Gaussian noise.

References

- [1] L. Albera, A. Ferréol, P. Comon, and P. Chevalier. Blind identification of overcomplete mixtures of sources (BIOME). *Linear algebra and its applications*, 391:3–30, 2004.
- [2] S. Arora, R. Ge, A. Moitra, and S. Sachdeva. Provable ICA with unknown Gaussian noise, with implications for Gaussian mixtures and autoencoders. In *NIPS*, pages 2384–2392, 2012.
- [3] J. Cardoso and A. Souloumiac. Blind beamforming for non-Gaussian signals. In *Radar and Signal Processing, IEE Proceedings F*, volume 140(6), pages 362–370. IET, 1993.
- [4] J.-F. Cardoso. Super-symmetric decomposition of the fourth-order cumulant tensor. Blind identification of more sources than sensors. In *Acoustics, Speech, and Signal Processing, 1991. ICASSP-91., 1991 International Conference on*, pages 3109–3112. IEEE, 1991.
- [5] J.-F. Cardoso and A. Souloumiac. Matlab JADE for real-valued data v 1.8. <http://perso.telecom-paristech.fr/~cardoso/Algo/Jade/jadeR.m>, 2005. [Online; accessed 8-May-2013].
- [6] P. Chevalier. Optimal separation of independent narrow-band sources: Concept and performance 1. *Signal Processing*, 73(12):27 – 47, 1999. ISSN 0165-1684.
- [7] P. Comon and C. Jutten, editors. *Handbook of Blind Source Separation*. Academic Press, 2010.
- [8] L. De Lathauwer, B. De Moor, and J. Vandewalle. Independent component analysis based on higher-order statistics only. In *Statistical Signal and Array Processing, 1996. Proceedings., 8th IEEE Signal Processing Workshop on*, pages 356–359. IEEE, 1996.
- [9] L. De Lathauwer, J. Castaing, and J. Cardoso. Fourth-order cumulant-based blind identification of underdetermined mixtures. *Signal Processing, IEEE Transactions on*, 55(6):2965–2973, June 2007. ISSN 1053-587X. doi: 10.1109/TSP.2007.893943.
- [10] H. Gävert, J. Hurri, J. Särelä, and A. Hyvärinen. Matlab FastICA v 2.5. <http://research.ics.aalto.fi/ica/fastica/code/dlcode.shtml>, 2005. [Online; accessed 1-May-2013].
- [11] N. Goyal, S. Vempala, and Y. Xiao. Fourier PCA. *CoRR*, abs/1306.5825, 2013.
- [12] A. Hyvärinen and E. Oja. Independent component analysis: Algorithms and applications. *Neural Networks*, 13(4-5):411–430, 2000.
- [13] A. Hyvärinen, J. Karhunen, and E. Oja. *Independent component analysis*. John Wiley & Sons, 2001.

- [14] M. Joho, H. Mathis, and R. H. Lambert. Overdetermined blind source separation: Using more sensors than source signals in a noisy mixture. In *Proc. International Conference on Independent Component Analysis and Blind Signal Separation. Helsinki, Finland*, pages 81–86, 2000.
- [15] Z. Koldovský and P. Tichavský. Methods of fair comparison of performance of linear ica techniques in presence of additive noise. In *Acoustics, Speech and Signal Processing, 2006. ICASSP 2006 Proceedings. 2006 IEEE International Conference on*, volume 5, pages V–V. IEEE, 2006.
- [16] Z. Koldovský and P. Tichavský. Asymptotic analysis of bias of fastica-based algorithms in presence of additive noise. Technical report, Technical report, 2007.
- [17] Z. Koldovský and P. Tichavský. Blind instantaneous noisy mixture separation with best interference-plus-noise rejection. In *Independent Component Analysis and Signal Separation*, pages 730–737. Springer, 2007.
- [18] S. Makino, T.-W. Lee, and H. Sawada. *Blind speech separation*. Springer, 2007.
- [19] B. D. Van Veen and K. M. Buckley. Beamforming: A versatile approach to spatial filtering. *IEEE assp magazine*, 5(2):4–24, 1988.
- [20] R. Vigário, J. Sarela, V. Jousmiki, M. Hamalainen, and E. Oja. Independent component approach to the analysis of EEG and MEG recordings. *Biomedical Engineering, IEEE Transactions on*, 47(5):589–593, 2000.
- [21] J. R. Voss, L. Rademacher, and M. Belkin. Fast algorithms for Gaussian noise invariant independent component analysis. In *Advances in Neural Information Processing Systems 26*, pages 2544–2552. 2013.
- [22] A. Yeredor. Blind source separation via the second characteristic function. *Signal Processing*, 80(5):897–902, 2000.
- [23] A. Yeredor. Non-orthogonal joint diagonalization in the least-squares sense with application in blind source separation. *Signal Processing, IEEE Transactions on*, 50(7):1545–1553, 2002.

A PEGI for Complex Signals

In Section 2, we showed how to perform gradient iteration ICA within a pseudo-Euclidean inner product space. In this appendix, we show how this PEGI algorithm can be extended to include complex valued signals. For clarity, we repeat the entire PEGI algorithmic construction from Section 2 with the necessary modifications to handle the complex setting.

Throughout this appendix, we assume a noisy ICA model $\mathbf{X} = \mathbf{A}\mathbf{S} + \boldsymbol{\eta}$ where $\boldsymbol{\eta}$ is an arbitrary Gaussian noise independent of \mathbf{S} . We also assume that $m \leq n$, that m is known, and that the columns of A are linearly dependent.

A.1 Fourth Cumulants of Complex Variables

The gradient iteration relies on the properties of cumulants. We will focus on the fourth cumulant, though similar constructions may be given using other even order cumulants of higher order. We will use two versions of the fourth cumulant which capture slightly different fourth order information. For a zero-mean random variable X , they may be defined as $\kappa_4(X) := \mathbb{E}[X^4] - 3\mathbb{E}[X^2]^2$ and $\kappa_4^*(X) := \mathbb{E}[X^2 X^{*2}] - 2\mathbb{E}[X X^*]^2 - \mathbb{E}[X^2]\mathbb{E}[X^{*2}]$. For real random variables, these two definitions are equivalent, and they come from two different conjugation schemes when constructing the fourth order cumulant [see 7, Chapter 5, Section 1.2]. However, in general, only κ_4^* is guaranteed to be real valued. The higher order cumulants have nice algebraic properties which make them useful for ICA:

1. (Independence) If X and Y are independent random variables, then $\kappa_4(X + Y) = \kappa_4(X) + \kappa_4(Y)$ and $\kappa_4^*(X + Y) = \kappa_4^*(X) + \kappa_4^*(Y)$.
2. (Homogeneity) If α is a scalar, then $\kappa_4(\alpha X) = \alpha^4 \kappa_4(X)$ and $\kappa_4^*(\alpha X) = |\alpha|^4 \kappa_4^*(X)$.
3. (Vanishing Gaussians) If X is normally distributed then $\kappa_4(X) = 0$ and $\kappa_4^*(X) = 0$.

In this appendix, we consider a noisy ICA model $\mathbf{X} = \mathbf{A}\mathbf{S} + \boldsymbol{\eta}$ where $\boldsymbol{\eta}$ is a $\mathbf{0}$ -mean (possibly complex) Gaussian and independent of \mathbf{S} . We consider the following functions defined on the unit sphere: $f(\mathbf{u}) := \kappa_4(\langle \mathbf{X}, \mathbf{u} \rangle)$ and $f_*(\mathbf{u}) := \kappa_4^*(\langle \mathbf{X}, \mathbf{u} \rangle)$. Then, expanding using the above properties we obtain:

$$\begin{aligned} f(\mathbf{u}) &= \kappa_4\left(\sum_{k=1}^m \langle A_k, \mathbf{u} \rangle S_k + \langle \mathbf{u}, \boldsymbol{\eta} \rangle\right) \\ &= \sum_{k=1}^m \langle A_k, \mathbf{u} \rangle^4 \kappa_4(S_k) \end{aligned}$$

Using similar reasoning, it can be seen that $f_*(\mathbf{u}) = \sum_{k=1}^m |\langle A_k, \mathbf{u} \rangle|^4 \kappa_4^*(S_k)$.

It turns out that some slightly non-standard notions of derivatives are most useful in constructing the gradient iteration in the complex setting. We use real derivatives for the gradient and we use the complex Hessian. In particular, expanding $u_k = x_k + iy_k$, we use the gradient operator $\nabla := \sum_{k=1}^n \mathbf{e}_k \frac{\partial}{\partial x_k}$. We make use of the operators $\partial u_k := \frac{1}{2}(\frac{\partial}{\partial x_k} - i\frac{\partial}{\partial y_k})$ and $\partial u_k^* := \frac{1}{2}(\frac{\partial}{\partial x_k} + i\frac{\partial}{\partial y_k})$ to define $\mathcal{H} := \sum_{j=1}^n \sum_{k=1}^n \mathbf{e}_k \mathbf{e}_j^T \partial u_k \partial u_j^*$. Applying this version of the Hessian is different than using real derivatives as in the gradient operation.

Taking derivatives, we obtain:

$$\nabla f(\mathbf{u}) = 4 \sum_{k=1}^m \langle A_k, \mathbf{u} \rangle^3 \kappa_4(S_k) A_k \quad (5)$$

$$\begin{aligned} \mathcal{H} f_*(\mathbf{u}) &= 4 \sum_{k=1}^m |\langle A_k, \mathbf{u} \rangle|^2 \kappa_4^*(S_k) A_k^* A_k^T \\ &= A^* D(\mathbf{u}) A^T \end{aligned} \quad (6)$$

where $D(\mathbf{u})$ is a diagonal matrix with entries $D(\mathbf{u})_{kk} = 4|\langle A_k, \mathbf{u} \rangle|^2 \kappa_4^*(S_k)$.

Algorithm 3 Recovers a column of A up to an unknown scaling factor when \mathbf{u}_0 is generically chosen.

Inputs: \mathbf{u}_0 (A unit vector), C , ∇f
 $k \leftarrow 1$
repeat
 $\mathbf{u}_k \leftarrow \nabla f(C^\dagger \mathbf{u}_{k-1}) / \|\nabla f(C^\dagger \mathbf{u}_{k-1})\|$
 $k \leftarrow k + 1$
until Convergence (up to a unit modulus factor)
return \mathbf{u}_k

A.2 Gradient Iteration in a Pseudo-Euclidean Space

We now demonstrate that the gradient iteration can be performed using a generalized notion of an inner product space in which the columns of A are orthogonal. The natural candidate for the “inner product space” would be to use $\langle \cdot, \cdot \rangle_*$ defined as $\langle \mathbf{u}, \mathbf{v} \rangle_* := \mathbf{u}^T (A^* A^T)^\dagger \mathbf{v}^*$. Clearly, $\langle A_i, A_j \rangle_* = \delta_{ij}$ gives the desired orthogonality property. However, there are two issues with this “inner product space”: First, it is only an inner product space when A is non-singular (invertible). This turns out not to be a major issue, and we will move forward largely ignoring this point. The second issue is more fundamental: We only have access to the matrix $A^* A^T$ in the noise free setting where $\text{cov}(\mathbf{X})^T = (A A^H)^T = A^* A^T$. In the noisy setting, we have access to matrices of the form $\mathcal{H}_f(\mathbf{u}) = A^* D(\mathbf{u}) A^T$ from equation (6) instead.

We consider a pseudo-Euclidean inner product defined as follows: Let $C = A^* D A^T$ where D is a diagonal matrix with non-zero diagonal entries, and define $\langle \cdot, \cdot \rangle_C$ by $\langle \mathbf{u}, \mathbf{v} \rangle_C = \mathbf{u}^T C^\dagger \mathbf{v}^*$. When D contains negative entries, this is not a proper inner product since C is not positive definite. In particular, $\langle A_k, A_k \rangle_C = A_k^T (A^* D A^T)^\dagger A_k^* = d_{kk}^{-1}$ may be negative. Nevertheless, when $k \neq j$, $\langle A_k, A_j \rangle_C = A_k^T (A^* D A^T)^\dagger A_j^* = 0$ gives that the columns of A are orthogonal in this space.

We define functions $\alpha_k : \mathbb{C}^n \rightarrow \mathbb{C}$ by $\alpha_k(\mathbf{u}) = (A^\dagger \mathbf{u})_k$ such that for any $\mathbf{u} \in \text{span}(A_1, \dots, A_m)$, then $\mathbf{u} = \sum_{i=1}^m \alpha_i(\mathbf{u}) A_i$ is the expansion of \mathbf{u} in its A_i basis. Continuing from equation (5), for any $\mathbf{u} \in S^{n-1}$ we see

$$\begin{aligned} \nabla f(C^\dagger \mathbf{u}) &= 4 \sum_{k=1}^n \langle A_k, C^\dagger \mathbf{u} \rangle^3 \kappa_4(S_k) A_k \\ &= 4 \sum_{k=1}^n \langle A_k, \mathbf{u} \rangle_C^3 \kappa_4(S_k) A_k \end{aligned}$$

is the gradient iteration recast in the $\langle \cdot, \cdot \rangle_C$ space. Expanding \mathbf{u} in its A_k basis, we obtain

$$\begin{aligned} \nabla f(C^\dagger \mathbf{u}) &= 4 \sum_{k=1}^m (\alpha_k(\mathbf{u}) \langle A_k, A_k \rangle_C)^3 \kappa_4(S_k) A_k \\ &= 4 \sum_{k=1}^m \alpha_k(\mathbf{u})^3 (d_{kk}^{-3} \kappa_4(S_k)) A_k, \end{aligned} \tag{7}$$

which is a power iteration in the unseen A_k coordinate system. As no assumptions are made upon the $\kappa_4(S_k)$ values, the d_{kk}^{-3} scalings which were not present in equation (5) cause no issues. Using this update, we obtain Algorithm 3, a fixed point method for recovering a single column of A up to an unknown scaling.

Before proceeding, we should clarify the notion of fixed point convergence in Algorithm 3. We say that the sequence $\{\mathbf{u}_k\}_{k=0}^\infty$ converges to \mathbf{v} up to a unit modulus factor if there exists a sequence of constants $\{c_k\}_{k=0}^\infty$ such that each $|c_k| = 1$ and $c_k \mathbf{u}_k \rightarrow \mathbf{v}$ as $k \rightarrow \infty$. We have the following convergence guarantee.

Theorem 5. *If \mathbf{u}_0 is chosen uniformly at random from S^{n-1} . Then with probability 1, there exists $\ell \in [m]$ such that the sequence $\{\mathbf{u}_k\}_{k=0}^\infty$ defined as in Algorithm 3 converges to a $A_\ell/\|A_\ell\|$ up to a unit modulus factor. Further, the rate of convergence is cubic.*

Due to space limitations, we omit the proof of Theorem 5. However, its proof is very similar that of an analogous result for the GI-ICA algorithm [21, Theorem 4].

In practice, we test near convergence by testing if we are still making significant progress. In particular, for some predefined $\epsilon > 0$, if there exists a unit modulus constant c_k such that $\|\mathbf{u}_k - c_k \mathbf{u}_{k-1}\| < \epsilon$, then we declare convergence achieved and return the result. We may determine c_k using the following fact.

Fact 6. Suppose that \mathbf{u} and \mathbf{v} are non-orthogonal unit modulus vectors. The expression $\|\mathbf{u} - e^{i\theta} \mathbf{v}\|$ is minimized by the choice of $\theta = \text{atan2}(\text{Im}(\langle \mathbf{u}, \mathbf{v} \rangle), \text{Re}(\langle \mathbf{u}, \mathbf{v} \rangle))$.

Letting $\theta = \text{atan2}(\text{Im}(\langle \mathbf{u}_k, \mathbf{u}_{k-1} \rangle), \text{Re}(\langle \mathbf{u}_k, \mathbf{u}_{k-1} \rangle))$, we exit the loop if $\|\mathbf{u}_k - e^{i\theta} \mathbf{u}_{k-1}\| < \epsilon$.

A.3 Full ICA Recovery Via the Pseudo-Euclidean GI-Update

We are able to recover a single column of A in noisy ICA. However, for full matrix recovery, we would like (given recovered columns $A_{\ell_1}, \dots, A_{\ell_j}$) to be able to recover a column A_k such that $k \notin \{\ell_1, \dots, \ell_j\}$ on demand.

The main idea behind the simultaneous recovery of all columns of A is two-fold. First, instead of just finding columns of A using Algorithm 3, we simultaneously find rows of A^\dagger . Then, using the recovered columns of A and rows of A^\dagger , we may project \mathbf{u} onto the orthogonal complement of the recovered columns of A within the $\langle \cdot, \cdot \rangle_C$ pseudo-Euclidean inner product space.

Recovering rows of A^\dagger . Suppose we have access to a column A_k (which may be achieved using Algorithm 3). Let $A_{k\cdot}^\dagger$ denote the k^{th} row of A^\dagger . Then, we note that $C^\dagger A_k^* = (A^* D A^T)^\dagger A_k^* = d_{kk}^{-1} (A^T)_{k\cdot}^\dagger = d_{kk}^{-1} (A_{k\cdot}^\dagger)^T$ recovers $A_{k\cdot}^\dagger$ up to an arbitrary, unknown constant d_{kk}^{-1} . However, the constant d_{kk}^{-1} may be recovered by noting that $\langle A_k, A_k \rangle_C = (C^\dagger A_k)^T A_k = d_{kk}^{-1}$. As such, we may estimate $A_{k\cdot}^\dagger$ as $[C^\dagger A_k / ((C^\dagger A_k)^T A_k)]^T$.

Enforcing Orthogonality During the GI Update. Given access to $\mathbf{u} = \sum_{k=1}^m \alpha_k(\mathbf{u}) A_k + P_{A^\perp} \mathbf{u}$, some recovered columns $A_{\ell_1}, \dots, A_{\ell_r}$, and corresponding rows of A^\dagger , we may zero out the components of \mathbf{u} corresponding to the recovered columns of A . Letting $\mathbf{u}' = \mathbf{u} - \sum_{j=1}^r A_{\ell_j} A_{\ell_j}^\dagger \mathbf{u}$, then $\mathbf{u}' = \sum_{k \in [m] \setminus \{\ell_1, \dots, \ell_r\}} \alpha_k(\mathbf{u}) A_k + P_{A^\perp} \mathbf{u}$. In particular, \mathbf{u}' is orthogonal (in the $\langle \cdot, \cdot \rangle_C$ space) to the previously recovered columns of A . This allows us to modify the non-orthogonal gradient iteration algorithm to recover a new column of A .

Using these ideas, we obtain the Algorithm 4 for recovery of the ICA mixing matrix. Within this Algorithm, step 6 enforces orthogonality with previously found columns of A , guaranteeing that convergence is to a new column of A .

Practical Construction of C We suggest the choice of $C = \frac{1}{4} \sum_{k=1}^n \mathcal{H}f_\star(\mathbf{e}_k)$, as it can be shown from equation (6) that $\sum_{k=1}^n \mathcal{H}f_\star(\mathbf{e}_k) = A^* D A^T$ with $d_{kk} = \|A_k\|^2 \kappa_4^*(S_k)$. This deterministically guarantees that each latent signal has a significant contribution to C .

Algorithm 4 Full ICA matrix recovery algorithm. Estimates and returns two matrices: (1) \tilde{A} is the recovered mixing matrix for the noisy ICA model $\mathbf{X} = \mathbf{A}\mathbf{S} + \boldsymbol{\eta}$, and (2) \tilde{B} is a running estimate of \tilde{A}^\dagger .

```

1: Inputs:  $C, \nabla f$ 
2:  $\tilde{A} \leftarrow 0, \tilde{B} \leftarrow 0$ 
3: for  $j \leftarrow 1$  to  $m$  do
4:   Draw  $\mathbf{u}$  uniformly at random from  $S^{n-1}$ .
5:   repeat
6:      $\mathbf{u} \leftarrow \mathbf{u} - \tilde{A}\tilde{B}\mathbf{u}$ 
7:      $\mathbf{u} \leftarrow \nabla f(C^\dagger^* \mathbf{u}) / \|\nabla f(C^\dagger^* \mathbf{u})\|$ .
8:   until Convergence (up to a unit modulus factor)
9:    $\tilde{A}_j \leftarrow \mathbf{u}$ 
10:   $\tilde{B}_j \leftarrow [C^\dagger A_j / ((C^\dagger A_j)^T A_j)]^T$ 
11: end for
12: return  $\tilde{A}, \tilde{B}$ 

```

B Proof of Proposition 2

Proof. This proof is based on the connection between two notions of optimality, minimum mean squared error and SINR. The mean squared error of the recovered signal $\hat{S}(\mathbf{b})$ from k^{th} latent signal is defined as $\text{MSE}_k(\mathbf{b}) := \mathbb{E}[|S_k - \hat{S}(\mathbf{b})|^2]$. It has been shown [14, equation 39] that B_{opt} jointly minimizes the mean squared errors of the recovered signals. In particular, if $\mathbf{b} = (B_{\text{opt}})_k$, then \mathbf{b} is a minimizer of $\text{MSE}_k(\mathbf{b})$.

We will first show that finding a matrix B which minimizes the mean squared error has the side effect of maximizing the magnitude of the Pearson correlations $\rho_{S_k, \hat{S}_k(B)}$ for each $k \in [m]$, where $\rho_{S_k, \hat{S}_k(B)} := \frac{\mathbb{E}[S_k \hat{S}_k^*(B)]}{\sigma_{S_k} \sigma_{\hat{S}_k(B)}}$. We will then demonstrate that if B is a maximizer of $|\rho_{S_k, \hat{S}_k(B)}|$, then B_k is a maximizer of SINR_k . These two facts imply the desired result. We will use the convention that $\rho_{S_k, \hat{S}_k(B)}$ is 0 if $\sigma_{\hat{S}_k(B)} = 0$.

We fix a $k \in [m]$. We have:

$$\begin{aligned} \text{MSE}_k(\mathbf{b}) &= \mathbb{E}[S_k S_k^* - 2 \text{Re}(S_k \hat{S}^*(\mathbf{b})) + \hat{S}(\mathbf{b}) \hat{S}^*(\mathbf{b})] \\ &= 1 - 2\sigma_{\hat{S}(\mathbf{b})} \text{Re}(\rho_{S_k, \hat{S}(\mathbf{b})}) + \sigma_{\hat{S}(\mathbf{b})}^2. \end{aligned}$$

Letting $\omega = \text{sgn}(\rho_{S_k, \hat{S}(\mathbf{b})})$, we obtain

$$\rho_{S_k, \hat{S}(\omega \mathbf{b})} = \frac{\mathbb{E}[S_k \hat{S}^*(\omega \mathbf{b})]}{\sigma_{S_k} \sigma_{\hat{S}(\omega \mathbf{b})}} = \omega^* \frac{\mathbb{E}[S_k \hat{S}^*(\mathbf{b})]}{\sigma_{S_k} \sigma_{\hat{S}(\mathbf{b})}} = |\rho_{S_k, \hat{S}(\mathbf{b})}|. \quad (8)$$

Further, $\text{MSE}_k(\omega \mathbf{b}) = 1 - 2\sigma_{\hat{S}(\mathbf{b})} |\rho_{S_k, \hat{S}(\mathbf{b})}| + \sigma_{\hat{S}(\mathbf{b})}^2 \leq \text{MSE}_k(\mathbf{b})$ with equality if and only if $\rho_{S_k, \hat{S}(\mathbf{b})}$ is real and non-negative. As such, all global minima of MSE_k are contained in the set $\mathcal{A} = \{\mathbf{b} \mid \rho_{S_k, \hat{S}(\mathbf{b})} \in [0, 1]\}$, and we may restrict our investigation to this set.

We define a function $g(x, y) := 1 - 2xy + y^2$ such that under the change of variable $x(\mathbf{b}) = \sigma_{\hat{S}(\mathbf{b})}$ and $y(\mathbf{b}) = \rho_{S_k, \hat{S}(\mathbf{b})}$, we obtain $\text{MSE}_k(\mathbf{b}) = g(x, y)$. Let $M = \max_{\mathbf{b} \in \mathcal{A}} \rho_{S_k, \hat{S}(\mathbf{b})}$ and let $y_0 \in [0, M]$ be fixed. Then, $\arg \min_{x \in \mathbb{R}} g(x, y_0) = y_0$ with the resulting value $g(y_0, y_0) = 1 - y_0^2$. As such, the minimum of $g(x, y)$ over the domain $\mathbb{R} \times [0, M]$ occurs when $x = y = M$. If $M = 0$, then the

choice of $\boldsymbol{\xi} = \mathbf{0}$ satisfies that $x(\boldsymbol{\xi}) = y(\boldsymbol{\xi}) = 0$, making $\text{MSE}_k(\boldsymbol{\xi}) = g(x, y)$ the global minimum of MSE_k . If $M \neq 0$, then we may choose $\boldsymbol{\xi}$ such that $y(\boldsymbol{\xi}) = \rho_{S_k, \hat{S}(\boldsymbol{\xi})} = M$. As $\sigma_{\hat{S}(\boldsymbol{\xi})} > 0$ must hold, it follows that there exists $\alpha \in (0, \infty)$ such that setting $\boldsymbol{\zeta} = \alpha\boldsymbol{\xi}$, we obtain $(\sigma_{\hat{S}(\boldsymbol{\zeta})} =) x(\boldsymbol{\zeta}) = y(\boldsymbol{\xi})$. Since $y(\mathbf{b}) = \rho_{S_k, \hat{S}(\mathbf{b})}$ is scale invariant, we obtain that $x(\boldsymbol{\zeta}) = y(\boldsymbol{\zeta}) = y(\boldsymbol{\xi}) = M$, making $\boldsymbol{\zeta}$ a global minimum of MSE_k . In both cases, it follows that if \mathbf{b} minimizes $\text{MSE}_k(\mathbf{b})$, then \mathbf{b} maximizes $\rho_{S_k, \hat{S}(\mathbf{b})}$ over \mathcal{A} .

From equation (8), we see that $\max_{\mathbf{b} \in \mathbb{C}^n} |\rho_{S_k, \hat{S}(\mathbf{b})}| = \max_{\mathbf{b} \in \mathcal{A}} \rho_{S_k, \hat{S}(\mathbf{b})}$. Thus if \mathbf{b} is a minimizer of $\text{MSE}_k(\omega\mathbf{b})$, then \mathbf{b} is also a maximizer of $|\rho_{S_k, \hat{S}(\mathbf{b})}|$ as claimed.

We now demonstrate that \mathbf{b} is a maximizer of $|\rho_{S_k, \hat{S}(\mathbf{b})}|$ if and only if it is also a maximizer of $\text{SINR}_k(\mathbf{b})$. Under the conventions that $\frac{x}{0} = +\infty$ when $x > 0$ and that $\frac{0}{0} = s\infty$ where $s = -1$ for maximization problems and $s = +1$ for minimization problems, the following problems have equivalent solution sets over choices of \mathbf{b} :

$$\begin{aligned} \max_{\mathbf{b}} \text{SINR}_k(\mathbf{b}) &\equiv \max_{\mathbf{b}} \frac{\mathbb{E}[|\mathbf{b}A_k S_k|^2]}{\text{var}(\hat{S}(\mathbf{b})) - \mathbb{E}[|\mathbf{b}A_k S_k|^2]} \equiv \max_{\mathbf{b}} \frac{|\mathbb{E}[S_k \hat{S}^*(\mathbf{b})]|^2}{\text{var}(\hat{S}(\mathbf{b})) - |\mathbb{E}[S_k \hat{S}^*(\mathbf{b})]|^2} \\ &\equiv \min_{\mathbf{b}} \frac{\text{var}(\hat{S}(\mathbf{b})) - |\mathbb{E}[S_k \hat{S}^*(\mathbf{b})]|^2}{|\mathbb{E}[S_k \hat{S}^*(\mathbf{b})]|^2} \equiv \min_{\mathbf{b}} \frac{\text{var}(\hat{S}(\mathbf{b}))}{|\mathbb{E}[S_k \hat{S}^*(\mathbf{b})]|^2} \\ &\equiv \max_{\mathbf{b}} \frac{|\mathbb{E}[S_k \hat{S}^*(\mathbf{b})]|^2}{\text{var}(\hat{S}(\mathbf{b}))} \equiv \max_{\mathbf{b}} |\rho_{S_k, \hat{S}(\mathbf{b})}|^2. \end{aligned}$$

In the above, the first equivalence is a rewriting of equation (4). To see the second equivalence, we note that $|\mathbb{E}[S_k \hat{S}^*(\mathbf{b})]|^2 = |\mathbb{E}[S_k(\mathbf{b}A\mathbf{S} + \mathbf{b}\boldsymbol{\eta})^*]|^2 = |\mathbf{b}A_k|^2$ using the independence of S_k from all other terms. Then, noting that $|\mathbf{b}A_k|^2 = \mathbb{E}[|\mathbf{b}A_k S_k|^2]$ gives the equivalence. The fourth equivalence is only changing the problem by the additive constant -1 . \square